arXiv:1603.06547v1 [math.LO] 21 Mar 2016

Constructive canonicity for lattice-based fixed point logics

Willem Conradie [1] Andrew Craig

Department of Pure and Applied Mathematics, University of Johannesburg, South Africa

Alessandra Palmigiano [2]

Department of Pure and Applied Mathematics, University of Johannesburg, South Africa

Faculty of Technology, Policy and Management, Delft University of Technology, the Netherlands

Zhiguang Zhao

Faculty of Technology, Policy and Management, Delft University of Technology, the Netherlands

Abstract

We prove the algorithmic canonicity of two classes of μ-inequalities in a constructive meta-theory of normal lattice expansions. This result simultaneously generalizes Conradie and Craig's canonicity for μ-inequalities based on a bi-intuitionistic bi-modal language, [4], and Conradie and Palmigiano's constructive canonicity for inductive inequalities [11] (restricted to normal lattice expansions to keep the page limit). Besides the greater generality, the unification of these strands smoothes the existing treatments for the canonicity of μ-formulas and inequalities. In particular, the rules of the algorithm ALBA used for this result have exactly the same formulation as those of [11], with no additional rule added specifically to handle the fixed point binders. Rather, fixed points are accounted for by certain restrictions on the application of the rules, concerning the order-theoretic properties of the term functions associated with the formulas to which the rules are applied.

Keywords: Constructive canonicity, modal mu-calculus, normal lattice expansions, Sahlqvist theory, unified correspondence.

1 Introduction

The present contribution lies at the crossroads of at least three active lines of research in nonclassical logics: the one investigating the semantic and proof-theoretic environ-

[1] The research of the first author has been made possible by the National Research Foundation of South Africa, Grant number 81309.

[2] The research of the third and fourth author has been made possible by the NWO Vidi grant 016.138.314, by the NWO Aspasia grant 015.008.054, and by a Delft Technology Fellowship awarded in 2013.

ment of fixed point expansions of logics algebraically captured by varieties of (distributive) lattice expansions [2,22,26,3,18]; the one investigating constructive canonicity for intuitionistic and substructural logics [20,27]; the one uniformly extending the state-of-the-art in Sahlqvist theory to families of nonclassical logics, and applying it to issues both semantic and proof-theoretic [8,12,10,6,15,25,24,13,14,11,21,23,17], known as 'unified correspondence'.

We prove the algorithmic canonicity of two classes of μ-inequalities in a constructive meta-theory of normal lattice expansions. This result simultaneously generalizes Conradie and Craig's canonicity results for μ-inequalities based on a bi-intuitionistic bi-modal language [4], and Conradie and Palmigiano's constructive canonicity for inductive inequalities [9] (restricted to normal lattice expansions to keep the page limit). Besides the greater generality, the unification of these strands smoothes the existing proofs for the canonicity of μ-formulas and inequalities. Specifically, the two canonicity results proven in [4], (namely, the tame and proper canonicity, cf. Section 3) fully generalize to the constructive setting and normal LEs. Remarkably, the rules of the algorithm ALBA used for this result have exactly the same formulation as those of [9], with no additional rule added specifically to handle the fixed point binders. Rather, fixed points are accounted for by certain restrictions on the application of the rules, concerning the order-theoretic properties of the term functions associated with the formulas to which the rules are applied.

Applications of these results include the formalization of common knowledge-type processes of social interaction. For instance, in ongoing work [7] we are exploring the formalization of processes giving rise to categorization systems (where categories arise, as in Formal Concept Analysis, as Galois-stable sets from a given polarity) agreed upon by a whole community.

Structure of the paper. In Section 2, we collect the needed preliminaries. In Section 4, we define the classes of mu-inequalities to which the canonicity results apply. In Section 3, we introduce and expand on the two notions of canonicity mentioned above. In Section 5, we outline the constructive and general lattice version of the algorithm ALBA which is the main tool for the two canonicity results (more details are given in the appendix). In Section 6, we state our main results. Due to space constraints we do not include proofs, but these may be found online in the an expanded version of the present paper [5].

2 Preliminaries

In this section we collect the needed preliminaries about the syntax and algebraic semantics of lattice-based mu-calculi.

2.1 The language of lattice expansions

In the present subsection, we introduce the propositional fragments of the lattice-based μ-languages we consider in this paper. We will make use of the following auxiliary definitions: an *order-type* over $n \in \mathbb{N}$[3] is an element $\varepsilon \in \{1, \partial\}^n$. For every order-type ε, we denote its *opposite* order-type by ε^∂, that is, $\varepsilon_i^\partial = 1$ iff $\varepsilon_i = \partial$ for every

[3] Throughout the paper, order-types will be typically associated with arrays of variables $\overline{p} := (p_1, \ldots, p_n)$.

$1 \leq i \leq n$. For any lattice $\mathbb{A}$, we let $\mathbb{A}^1 := \mathbb{A}$ and $\mathbb{A}^\partial$ be the dual lattice, that is, the lattice associated with the converse partial order of $\mathbb{A}$. For any order-type ε, we let $\mathbb{A}^\varepsilon := \Pi_{i=1}^n \mathbb{A}^{\varepsilon_i}$. Sometimes we will write $\top^1$ and $\top^\partial$ for $\top$ and $\bot$ respectively. Similarly, we will write $\bot^1$ and $\bot^\partial$ for $\bot$ and $\top$ respectively. For both order-types and tuples of variables, we will use the symbol $\oplus$ to denote concatenation.

The LE-language $\mathcal{L}(\mathcal{F}, \mathcal{G})$ (abbreviated as $\mathcal{L}$ when no confusion arises) takes as parameters: 1) a denumerable set PROP of proposition letters, elements of which are denoted p, q, r, possibly with indices; 2) disjoint sets of connectives $\mathcal{F}$ and $\mathcal{G}$. Each $f \in \mathcal{F}$ (resp. $g \in \mathcal{G}$) has arity $n_f \in \mathbb{N}$ (resp. $n_g \in \mathbb{N}$), and is associated with some order-type ε_f over n_f (resp. ε_g over n_g).The terms (formulas) of $\mathcal{L}$ are defined recursively as follows:

$$\varphi ::= p \mid \bot \mid \top \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid f(\overline{\varphi}) \mid g(\overline{\varphi})$$

where $p \in \mathsf{PROP}$, $f \in \mathcal{F}$, $g \in \mathcal{G}$. Terms in $\mathcal{L}$ will be denoted either by s, t, or by lowercase Greek letters such as φ, ψ, γ etc.

2.2 Lattice expansions, and their canonical extensions

The following definition captures the algebraic setting of the present paper, namely the normal lattice expansions of [11].

Definition 2.1 For any LE-signature $\mathcal{L} = \mathcal{L}(\mathcal{F}, \mathcal{G})$, a *lattice expansion* (LE) is a tuple $\mathbb{A} = (L, \mathcal{F}^\mathbb{A}, \mathcal{G}^\mathbb{A})$ such that L is a bounded lattice, $\mathcal{F}^\mathbb{A} = \{f^\mathbb{A} \mid f \in \mathcal{F}\}$ and $\mathcal{G}^\mathbb{A} = \{g^\mathbb{A} \mid g \in \mathcal{G}\}$, such that every $f^\mathbb{A} \in \mathcal{F}^\mathbb{A}$ (resp. $g^\mathbb{A} \in \mathcal{G}^\mathbb{A}$) is an n_f-ary (resp. n_g-ary) operation on $\mathbb{A}$. An LE $\mathbb{A}$ is *normal* if every $f^\mathbb{A} \in \mathcal{F}^\mathbb{A}$ (resp. $g^\mathbb{A} \in \mathcal{G}^\mathbb{A}$) preserves finite (hence also empty) joins (resp. meets) in each coordinate with $\varepsilon_f(i) = 1$ (resp. $\varepsilon_g(i) = 1$) and reverses finite (hence also empty) meets (resp. joins) in each coordinate with $\varepsilon_f(i) = \partial$ (resp. $\varepsilon_g(i) = \partial$). Let $\mathbb{LE}$ be the class of LEs. Sometimes we will refer to certain LEs as $\mathcal{L}$-algebras to emphasize the signature. In the remainder of the paper, we will abuse notation and write e.g. f for $f^\mathbb{A}$ when this causes no confusion. From now on, normal lattice expansions will be abbreviated as LEs.

Each language $\mathcal{L}$ is interpreted in the appropriate class of LEs. In particular, for every LE $\mathbb{A}$, each operation $f^\mathbb{A} \in \mathcal{F}^\mathbb{A}$ (resp. $g^\mathbb{A} \in \mathcal{G}^\mathbb{A}$) is finitely join-preserving (resp. meet-preserving) in each coordinate when regarded as a map $f^\mathbb{A} : \mathbb{A}^{\varepsilon_f} \to \mathbb{A}$ (resp. $g^\mathbb{A} : \mathbb{A}^{\varepsilon_g} \to \mathbb{A}$).

2.3 The 'tense' language $\mathcal{L}^t$

Any language $\mathcal{L} = \mathcal{L}(\mathcal{F}, \mathcal{G})$ can be associated with the language $\mathcal{L}^t = \mathcal{L}_{\mathrm{LE}}(\mathcal{F}^t, \mathcal{G}^t)$, where $\mathcal{F}^t \supseteq \mathcal{F}$ and $\mathcal{G}^t \supseteq \mathcal{G}$ are obtained by adding:

(i) for $f \in \mathcal{F}$ and $1 \leq i \leq n_f$, the n_f-ary connective $f_i^\sharp$, the intended interpretation of which is the right residual of f in its ith coordinate if $\varepsilon_f(i) = 1$ (resp. its Galois-adjoint if $\varepsilon_f(i) = \partial$);

(ii) for $g \in \mathcal{G}$ and $1 \leq i \leq n_g$, the n_g-ary connective $g_i^\flat$, the intended interpretation of which is the left residual of g in its ith coordinate if $\varepsilon_g(i) = 1$ (resp. its Galois-adjoint if $\varepsilon_g(i) = \partial$).

We stipulate that $f_i^\sharp \in \mathcal{G}^t$ if $\varepsilon_f(i) = 1$, and $f_i^\sharp \in \mathcal{F}^t$ if $\varepsilon_f(i) = \partial$. Dually, $g_i^\flat \in \mathcal{F}^t$ if $\varepsilon_g(i) = 1$, and $g_i^\flat \in \mathcal{G}^t$ if $\varepsilon_g(i) = \partial$. The order-type assigned to the additional connectives is predicated on the order-type of their intended interpretations. That is, for any $f \in \mathcal{F}$ and $g \in \mathcal{G}$,

(i) if $\varepsilon_f(i) = 1$, then $\varepsilon_{f_i^\sharp}(i) = 1$ and $\varepsilon_{f_i^\sharp}(j) = (\varepsilon_f(j))^\partial$ for any $j \neq i$.

(ii) if $\varepsilon_f(i) = \partial$, then $\varepsilon_{f_i^\sharp}(i) = \partial$ and $\varepsilon_{f_i^\sharp}(j) = \varepsilon_f(j)$ for any $j \neq i$.

(iii) if $\varepsilon_g(i) = 1$, then $\varepsilon_{g_i^\flat}(i) = 1$ and $\varepsilon_{g_i^\flat}(j) = (\varepsilon_g(j))^\partial$ for any $j \neq i$.

(iv) if $\varepsilon_g(i) = \partial$, then $\varepsilon_{g_i^\flat}(i) = \partial$ and $\varepsilon_{g_i^\flat}(j) = \varepsilon_g(j)$ for any $j \neq i$.

The algebraic semantics of $\mathcal{L}^t := \mathcal{L}(!^t, \mathcal{G}^t)$ is given by the class of 'tense' $\mathcal{L}^t$-algebras, defined as those $\mathcal{L}^t$-algebras $\mathbb{A} = (L, (\mathcal{F}^t)^{\mathbb{A}}, (\mathcal{G}^t)^{\mathbb{A}})$ such that

(i) for every $f \in \mathcal{F}$ s.t. $n_f \geq 1$, all $a_1, \ldots, a_{n_f}, b \in L$, and each $1 \leq i \leq n_f$,
- if $\varepsilon_f(i) = 1$, then $f(a_1, \ldots, a_i, \ldots a_{n_f}) \leq b$ iff $a_i \leq f_i^\sharp(a_1, \ldots, b, \ldots, a_{n_f})$;
- if $\varepsilon_f(i) = \partial$, then $f(a_1, \ldots, a_i, \ldots a_{n_f}) \leq b$ iff $a_i \leq^\partial f_i^\sharp(a_1, \ldots, b, \ldots, a_{n_f})$.

(ii) for every $g \in \mathcal{G}$ s.t. $n_g \geq 1$, any $a_1, \ldots, a_{n_g}, b \in L$, and each $1 \leq i \leq n_g$,
- if $\varepsilon_g(i) = 1$, then $b \leq g(a_1, \ldots, a_i, \ldots a_{n_g})$ iff $g_i^\flat(a_1, \ldots, b, \ldots, a_{n_g}) \leq a_i$.
- if $\varepsilon_g(i) = \partial$, then $b \leq g(a_1, \ldots, a_i, \ldots a_{n_g})$ iff $g_i^\flat(a_1, \ldots, b, \ldots, a_{n_g}) \leq^\partial a_i$.

2.4 Canonical extensions, constructively

Canonical extensions provide a purely algebraic encoding of Stone-type dualities, and indeed, the existence of the canonical extensions of the best-known varieties of LEs can be proven via preexisting dualities. However, alternative, purely algebraic constructions are available, such as those of [19,16]. These constructions are in fact more general, in that their definition does not rely on principles such as Zorn's lemma. In what follows we will recall them relative to the setting of LEs introduced above.

Definition 2.2 Let $\mathbb{A}$ be a (bounded) sublattice of a complete lattice $\mathbb{A}'$.

(i) $\mathbb{A}$ is *dense* in $\mathbb{A}'$ if every element of $\mathbb{A}'$ can be expressed both as a join of meets and as a meet of joins of elements from $\mathbb{A}$.

(ii) $\mathbb{A}$ is *compact* in $\mathbb{A}'$ if, for all $S, T \subseteq \mathbb{A}'$, if $\bigvee S \leq \bigwedge T$ then $\bigvee S' \leq \bigwedge T'$ for some finite $S' \subseteq S$ and $T' \subseteq T$.

(iii) The *canonical extension* of a lattice $\mathbb{A}$ is a complete lattice $\mathbb{A}^\delta$ containing $\mathbb{A}$ as a dense and compact sublattice.

Let $K(\mathbb{A}^\delta)$ and $O(\mathbb{A}^\delta)$ denote the meet-closure and the join-closure of $\mathbb{A}$ in $\mathbb{A}^\delta$ respectively. The elements of $K(\mathbb{A}^\delta)$ are referred to as *closed* elements, and elements of $O(\mathbb{A}^\delta)$ as *open* elements. The canonical extension of a bounded lattice $\mathbb{A}$ exists and is unique up to any isomorphism fixing $\mathbb{A}$ (cf. [19, Propositions 2.6 and 2.7]). In meta-theoretic settings in which Zorn's lemma is available, the canonical extension of a lattice $\mathbb{A}$ is a *perfect* lattice. That is, in addition to being complete, is both completely join-generated by the set $J^\infty(\mathbb{A})$ of the completely join-irreducible elements of $\mathbb{A}$, and completely meet-generated by the set $M^\infty(\mathbb{A})$ of the completely meet-irreducible elements of $\mathbb{A}$. In our present, constructive setting, canonical extensions might not be

perfect.

The canonical extension of an LE $\mathbb{A}$ will be defined as a suitable expansion of the canonical extension of the underlying lattice of $\mathbb{A}$. Before turning to this definition, recall that taking order-duals interchanges closed and open elements: $K((\mathbb{A}^\delta)^\partial) = O(\mathbb{A}^\delta)$ and $O((\mathbb{A}^\delta)^\partial) = K(\mathbb{A}^\delta)$; similarly, $K((\mathbb{A}^n)^\delta) = K(\mathbb{A}^\delta)^n$, and $O((\mathbb{A}^n)^\delta) = O(\mathbb{A}^\delta)^n$. Hence, $K((\mathbb{A}^\delta)^\varepsilon) = \prod_i K(\mathbb{A}^\delta)^{\varepsilon(i)}$ and $O((\mathbb{A}^\delta)^\varepsilon) = \prod_i O(\mathbb{A}^\delta)^{\varepsilon(i)}$ for every LE $\mathbb{A}$ and every order-type ε on any $n \in \mathbb{N}$, where

$$K(\mathbb{A}^\delta)^{\varepsilon(i)} := \begin{cases} K(\mathbb{A}^\delta) & \text{if } \varepsilon(i) = 1 \\ O(\mathbb{A}^\delta) & \text{if } \varepsilon(i) = \partial \end{cases} \qquad O(\mathbb{A}^\delta)^{\varepsilon(i)} := \begin{cases} O(\mathbb{A}^\delta) & \text{if } \varepsilon(i) = 1 \\ K(\mathbb{A}^\delta) & \text{if } \varepsilon(i) = \partial. \end{cases}$$

As a consequence of these observations, taking the canonical extension of a lattice commutes with taking order duals and products, namely: $(\mathbb{A}^\partial)^\delta = (\mathbb{A}^\delta)^\partial$ and $(\mathbb{A}_1 \times \mathbb{A}_2)^\delta = \mathbb{A}_1^\delta \times \mathbb{A}_2^\delta$ (cf. [16, Theorem 2.8]). Hence, $(\mathbb{A}^\partial)^\delta$ can be identified with $(\mathbb{A}^\delta)^\partial$, $(\mathbb{A}^n)^\delta$ with $(\mathbb{A}^\delta)^n$, and $(\mathbb{A}^\varepsilon)^\delta$ with $(\mathbb{A}^\delta)^\varepsilon$ for any order-type ε. Thanks to these identifications, in order to extend operations of any arity and which are monotone or antitone in each coordinate from a lattice $\mathbb{A}$ to its canonical extension, treating the case of *monotone* and *unary* operations suffices:

Definition 2.3 For every unary, order-preserving operation $f : \mathbb{A} \to \mathbb{A}$, the σ- and π-*extension* of f are defined as follows:

$$f^\sigma(u) = \bigvee\{\bigwedge\{f(a) : k \leq a \in \mathbb{A}\} : u \geq k \in K(\mathbb{A}^\delta)\}$$
$$f^\pi(u) = \bigwedge\{\bigvee\{f(a) : o \geq a \in \mathbb{A}\} : u \leq o \in O(\mathbb{A}^\delta)\}.$$

It is easy to see that the σ- and π-extensions of ε-monotone maps are ε-monotone. Moreover, the σ-extension of a map which sends finite (resp. finite nonempty) joins or meets in the domain to finite (resp. finite nonempty) joins in the codomain sends *arbitrary* (resp. *arbitrary* nonempty) joins or meets in the domain to *arbitrary* (resp. *arbitrary* nonempty) joins in the codomain. Dually, the π-extension of a map which sends finite (resp. finite nonempty) joins or meets in the domain to finite (resp. finite nonempty) meets in the codomain sends *arbitrary* (resp. *arbitrary* nonempty) joins or meets in the domain to *arbitrary* (resp. *arbitrary* nonempty) meets in the codomain (cf. [19, Lemma 4.6]; notice that the proof given there holds in a constructive metatheory). Therefore, depending on the properties of the original operation, it is more convenient to use one or the other extension. This justifies the following:

Definition 2.4 The canonical extension of an $\mathcal{L}$-algebra $\mathbb{A} = (L, \mathcal{F}^\mathbb{A}, \mathcal{G}^\mathbb{A})$ is the $\mathcal{L}$-algebra $\mathbb{A}^\delta := (L^\delta, \mathcal{F}^{\mathbb{A}^\delta}, \mathcal{G}^{\mathbb{A}^\delta})$ such that $f^{\mathbb{A}^\delta}$ and $g^{\mathbb{A}^\delta}$ are defined as the σ-extension of $f^\mathbb{A}$ and as the π-extension of $g^\mathbb{A}$ respectively, for all $f \in \mathcal{F}$ and $g \in \mathcal{G}$.

The canonical extension of an LE $\mathbb{A}$ is a *quasi-perfect* LE:

Definition 2.5 An LE $\mathbb{A} = (L, \mathcal{F}^\mathbb{A}, \mathcal{G}^\mathbb{A})$ is quasi-perfect if L is a complete lattice, and the infinitary versions of the distribution laws defining normality are satisfied for each $f \in \mathcal{F}$ and $g \in \mathcal{G}$.

2.5 Adding the fixed point operators

In this subsection, we describe two ways of extending any LE-language by adding fixed point operators. The distinction between the two extensions will become clear

when we define their interpretations on LEs. In what follows, FVAR, PHVAR are disjoint sets of *fixed point variables* and *placeholder variables*, respectively.

For any LE-language $\mathcal{L}$, let $\mathcal{L}_1$ be the set of terms which extends $\mathcal{L}$ by allowing terms $\mu x.t(x)$ and $\nu x.t(x)$ where $t \in \mathcal{L}_1$, $x \in$ FVAR and $t(x)$ is positive in x. The second extension is denoted $\mathcal{L}_2$ and extends $\mathcal{L}$ by allowing construction of the terms $\mu_2 x.t(x)$ and $\nu_2 x.t(x)$ where $t \in \mathcal{L}_2$, $x \in$ FVAR and $t(x)$ is positive in x.

Terms in $\mathcal{L}$ are interpreted in LEs as described above. If $t(x_1, x_2, \ldots, x_n) \in \mathcal{L}_1$ and $a_1, \ldots, a_{n-1} \in \mathbb{A}$, then

$$\mu x.t(x, a_1, \ldots, a_{n-1}) := \bigwedge\{ a \in \mathbb{A} \mid t(a, a_1, \ldots, a_{n-1}) \leq a \}$$

if this meet exists, otherwise $\mu x.t(x, a_1, \ldots, a_{n-1})$ is undefined. Similarly,

$$\nu x.t(x, a_1, \ldots, a_{n-1}) := \bigvee\{ a \in \mathbb{A} \mid a \leq t(a, a_1, \ldots, a_{n-1}) \}$$

if this join exists, otherwise $\nu x.t(x, a_1, \ldots, a_{n-1})$ is undefined. For each ordinal α we define $t^\alpha(\bot, a_2, \ldots, a_n)$ as follows:

$$\begin{array}{lll}
t^0(\bot, a_1, \ldots, a_{n-1}) = \bot, & t^{\alpha+1}(\bot, a_1, \ldots, a_{n-1}) = t(t^\alpha(\bot, a_1, \ldots, a_{n-1}), a_1, \ldots, a_{n-1}), & \\
t^\lambda(\bot, a_1, \ldots, a_{n-1}) = \bigvee_{\alpha<\lambda} t^\alpha(\bot, a_1, \ldots, a_{n-1}) & & \text{for limit ordinals } \lambda; \\
t_0(\top, a_1, \ldots, a_{n-1}) = \top, & t_{\alpha+1}(\top, a_1, \ldots, a_{n-1}) = t(t_\alpha(\top, a_1, \ldots, a_{n-1}), a_1, \ldots, a_{n-1}), & \\
t_\lambda(\top, a_1, \ldots, a_{n-1}) = \bigwedge_{\alpha<\lambda} t_\alpha(\top, a_1, \ldots, a_{n-1}) & & \text{for limit ordinals } \lambda.
\end{array}$$

If $t(x_1, \ldots, x_n) \in \mathcal{L}_2$, then we let

$$\mu_2 x.t(x, a_1, \ldots, a_{n-1}) := \bigvee_{\alpha \geq 0} t^\alpha(\bot, a_1, \ldots, a_{n-1})$$
$$\nu_2 x.t(x, a_1, \ldots, a_{n-1}) := \bigwedge_{\alpha \geq 0} t_\alpha(\top, a_1, \ldots, a_{n-1})$$

if this join and meet exist, otherwise are undefined.

A lattice expansion $\mathbb{A}$ is *of the first kind* (*of the second kind*) if $t^{\mathbb{A}}(a_1, \ldots, a_n)$ is defined for all $a_1, \ldots, a_n \in \mathbb{A}$ and all $t \in \mathcal{L}_1$ ($t \in \mathcal{L}_2$). Henceforth we will refer to these algebras as *mu-algebras of the first kind* (*of the second kind*). When restricted to the Boolean case, our mu-algebras of the first kind are essentially the modal mu-algebras defined in [3, Definition 2.2] and [1, Definition 5.1]. Every mu-algebra of the second kind is a mu-algebra of the first kind. (cf. [1, Proposition 2.4] and [4, Lemma 2.2]. These proofs straightforwardly extend to the setting of general LEs). Hence, the interpretation of the two types of fixed point binders on mu-algebras of the second kind will agree. That is, $\mu X.\varphi(X) = \mu_2 X.\varphi(X)$ and $\nu X.\psi(X) = \nu_2 X.\psi(X)$ in mu-algebras of the second kind.

The final sets of terms, $\mathcal{L}_*$ (resp. $\mathcal{L}^t_*$), are obtained as extensions of $\mathcal{L}$ (resp. $\mathcal{L}^t$) by allowing $\mu^* x.s(x)$ and $\nu^* x.s(x)$ whenever $s \in \mathcal{L}_*$ (resp. $s \in \mathcal{L}^t_*$) and is positive in x. Terms in $\mathcal{L}_*$ and $\mathcal{L}^t_*$ are only interpreted in the constructive canonical extensions $\mathbb{A}^\delta$ of lattice expansions $\mathbb{A}$. If $s(x_1, x_2, \ldots, x_n) \in \mathcal{L}_* \cup \mathcal{L}^t_*$ and $a_1, \ldots, a_{n-1} \in \mathbb{A}^\delta$, then $\mu^* x_1.s(x_1, a_1, \ldots, a_{n-1}) := \bigwedge\{ a \in \mathbb{A} \mid s(a, a_1, \ldots, a_{n-1}) \leq a \}$ and $\nu^* x_1.s(x_1, a_2, \ldots, a_n) := \bigvee\{ a \in \mathbb{A} \mid a \leq s(a, a_1, \ldots, a_{n-1}) \}$. As the canonical extension $\mathbb{A}^\delta$ is a complete lattice, the interpretation of $\mu^* x.t(x)$ or $\nu^* x.t(x)$ is always defined. For any term $\varphi \in \mathcal{L}^t_1$ we let φ^* denote the $\mathcal{L}^t_*$ term obtained from φ by replacing all occurrences of μ and ν with μ^* and ν^*, respectively. The main feature of the μ^* and ν^* binders is that their interpretation does not change from $\mathbb{A}$ to $\mathbb{A}^\delta$.

2.6 The language of constructive ALBA for LEs

The expanded language manipulated by the algorithm μ^*-ALBA (cf. section 5) includes the $\mathcal{L}^t$-connectives, as well as a denumerably infinite set of sorted variables NOM called *nominals*, a denumerably infinite set of sorted variables CO-NOM, called *co-nominals*, and fixed point binders (depending on the language). We let $\mathbf{i}, \mathbf{j}$ denote nominals, and $\mathbf{m}, \mathbf{n}$ denote co-nominals. While in the non-constructive setting nominals and co-nominals range over the completely join-irreducible and the completely meet-irreducible elements of perfect LEs, respectively, in the present, constructive setting, nominals and co-nominals will be interpreted as elements of $K(\mathbb{A}^\delta)$ and $O(\mathbb{A}^\delta)$, respectively.

Formulas in the extended language $\mathcal{L}_1^+$ are defined by the following recursion:

$$\varphi ::= \bot \mid \top \mid p \mid X \mid \mathbf{j} \mid \mathbf{m} \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid f(\overline{\varphi}) \mid g(\overline{\varphi}) \mid \mu X.\varphi(X) \mid \nu X.\varphi(X)$$

where $p \in$ PROP, $X \in$ FVAR, $\mathbf{j} \in$ NOM, $\mathbf{m} \in$ CNOM, $f \in \mathcal{F}^t$, $g \in \mathcal{G}^t$, and φ is positive in X in $\mu X.\varphi(X)$ and $\nu X.\varphi(X)$. Formulas in $\mathcal{L}_2^+$ (resp. $\mathcal{L}_*^+$) are defined by replacing the fixed point operators μ, ν with μ_2, ν_2 (resp. μ^*, ν^*).

Placeholder variables from PHVAR, denoted x, y, z, will be used as generic variables which can take on the roles of propositional and fixed point variables. They will also be used to enhance the clarity of the exposition when dealing with substitution instances of formulas. Let τ be an order-type over n. An $\mathcal{L}_1^+$ (resp. $\mathcal{L}_2^+$, $\mathcal{L}_*^+$) formula is *pure* if it contains no ordinary (propositional) variables but only, possibly, nominals and co-nominals, and is an $\mathcal{L}_1^+$ (resp. $\mathcal{L}_2^+$, $\mathcal{L}_*^+$)-*sentence* if it contains no free fixed point variables.

A *quasi-inequality* of $\mathcal{L}_1^+$ (resp. $\mathcal{L}_2^+$, $\mathcal{L}_*^+$) is an expression of the form $\varphi_1 \leq \psi_1 \,\&\, \cdots \,\&\, \varphi_n \leq \psi_n \Rightarrow \varphi \leq \psi$ where the φ_i, ψ_i, φ and ψ are formulas of $\mathcal{L}_1^+$ (resp. $\mathcal{L}_2^+$, $\mathcal{L}_*^+$).

Quasi-inequalities, assignments, validity. Constructive canonical extensions of $\mathcal{L}$-algebras can be naturally endowed with the structure of $\mathcal{L}^+$-algebras (cf. Section 2.6 [9]). Building on this fact, we can use constructive canonical extensions of $\mathcal{L}$-algebras as a semantic environment for the language $\mathcal{L}^+$ as follows. For any $\mathcal{L}$-algebra $\mathbb{A}$, an *assignment* on $\mathbb{A}$ sends propositional variables to elements of $\mathbb{A}$ and is extended to formulas of $\mathcal{L}_1$, $\mathcal{L}_2$ and $\mathcal{L}_*$ in the usual way, where these are defined. An *assignment* on $\mathbb{A}^\delta$ is a map v : PROP$\cup$NOM$\cup$CO-NOM $\to \mathbb{A}^\delta$ sending propositional variables to elements of $\mathbb{A}^\delta$, nominals to $K(\mathbb{A}^\delta)$ and co-nominals to $O(\mathbb{A}^\delta)$ and extends to all formulas of $\mathcal{L}_*^+$, $\mathcal{L}_1^+$ and $\mathcal{L}_2^+$. An *admissible assignment* on $\mathbb{A}^\delta$ is an assignment which takes all propositional variables to elements of $\mathbb{A}$. An $\mathcal{L}^+$-inequality $\alpha \leq \beta$ is *admissibly valid* on $\mathbb{A}$, denoted $\mathbb{A}^\delta \models_{\mathbb{A}} \alpha \leq \beta$, if it holds under all admissible assignments. A quasi-inequality $\varphi_1 \leq \psi_1 \,\&\, \cdots \,\&\, \varphi_1 \leq \psi_1 \Rightarrow \varphi \leq \psi$ is satisfied under an assignment V in an algebra $\mathbb{A}$ of the appropriate sort, written $\mathbb{A}, V \models \varphi_1 \leq \psi_1 \,\&\, \cdots \,\&\, \varphi_n \leq \psi_n \Rightarrow \varphi \leq \psi$ if $\mathbb{A}, V \not\models \varphi_i \leq \psi_i$ for some $1 \leq i \leq n$ or $\mathbb{A}, v \models \varphi \leq \psi$. A quasi-inequality is (admissibly) valid in an algebra if it is satisfied by every (admissible) assignment.

Signed generation trees. For any formula/term φ in $\mathcal{L}_1^+$ and $\mathcal{L}_*^+$, we assign two *signed generation trees* $+\varphi$ and $-\varphi$. The generation tree is constructed as usual, beginning at the root with the main connective and then branching out into n-nodes at each n-ary connective. Each leaf is either a propositional variable, a fixed point variable, or

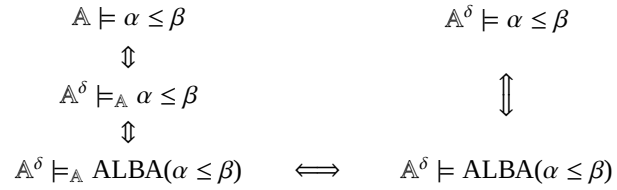
$\mathbb{A} \models \alpha \leq \beta$
$\mathbb{A}^\delta \models \alpha \leq \beta$
$\mathbb{A}^\delta \models_{\mathbb{A}} \alpha \leq \beta$
$\mathbb{A}^\delta \models_{\mathbb{A}} \mathrm{ALBA}(\alpha \leq \beta)$
$\iff$
$\mathbb{A}^\delta \models \mathrm{ALBA}(\alpha \leq \beta)$

Fig. 1. The U-shaped argument for canonicity of inequalities interpreted on a LE $\mathbb{A}$.

a constant. Each node is signed as follows:

- the root node of $+\varphi$ is signed $+$ and the root node of $-\varphi$ is signed $-$;
- if a node is $\vee, \wedge$ assign the same sign to its successor nodes;
- if a node is $h \in \mathcal{F}^t \cup \mathcal{G}^t$, assign the same (resp. the opposite) sign to every node corresponding to a coordinate i such that $\varepsilon_h(i) = 1$ (resp. $\varepsilon_h(i) = \partial$);
- if a node is $\mu x.\varphi(x)$, $\mu^* x.\varphi(x)$, $\nu x.\varphi(x)$ or $\nu^* x.\varphi(x)$ (with every free occurrence of x in the positive generation tree of φ labelled positively) then assign the same sign to the successor node.

A node in a signed generation tree is *positive* if it is signed "+" and *negative* if it is signed "−".

3 Two kinds of canonicity, constructively

In this section we give a brief conceptual and methodological overview of the main results of this paper.

The arguments for canonicity have the "U-shaped" format typical of the unified correspondence paradigm (see [8]) as generically illustrated in Figure 1. Going down the left-hand arm of the diagram, the first bi-implication is justified by the fact that validity in $\mathbb{A}$ is the same as admissible validity in $\mathbb{A}^\delta$, provided we make some stipulations regarding the interpretation of fixed point binders. The extended language $\mathcal{L}^+_*$ can be interpreted in $\mathbb{A}^\delta$ with nominals and co-nominals running over closed and open elements, respectively.

The inequality is now equivalently transformed into a set of pure (quasi-) inequalities, denoted $\mathrm{ALBA}(\alpha \leq \beta)$ in Figure 1. This is done by means of a calculus of rewrite rules encapsulated in the constructive μ^*-ALBA, which we present in Section 5. The fact that admissible and ordinary validity coincide for pure inequalities allows us to traverse the the bi-implication forming the base of the "U".

We proceed up the right-hand arm of the "U" by reversing the rewrite rules applied when going down the left-hand side. The equivalences are justified by the fact that these rules preserve validity on quasi-perfect algebras (cf. Definition 2.5).

If φ is a formula without fixed point binders, then the term function $\varphi^{\mathbb{A}^\delta}$ *extends* the term function $\varphi^{\mathbb{A}}$, i.e., they agree on arguments from $\mathbb{A}$. This is usually of crucial importance in proving that an equation is canonical. As soon as we add fixed point binders, this extension property fails. Indeed, $(\varphi(X))^{\mathbb{A}^\delta}$ can have more pre-fixed points in $\mathbb{A}^\delta$ than $(\varphi(X))^{\mathbb{A}}$ has in $\mathbb{A}$, and so $(\mu X.\varphi(X))^{\mathbb{A}^\delta}$ would generally be smaller than

$(\mu X.\varphi(X))^{\mathbb{A}}$. This phenomenon creates significant obstacles for standard canonicity arguments, and is not an additional difficulty posed by the constructive environment.

One way around these difficulties, adopted by Bezhanishvili and Hodkinson [3], is to require that only pre-fixed points from the smaller algebra $\mathbb{A}$ are used in calculating $(\mu X.\varphi(X))^{\mathbb{A}^\delta}$. This is tantamount to interpreting $\mu X.\varphi(X)$ in $\mathbb{A}^\delta$ as $(\mu^* X.\varphi(X))^{\mathbb{A}^\delta}$.Thus, in [4], two different notions of canonicity for the mu-calculus were considered, and their two ensuing canonicity results were shown for certain classes of mu-inequalities. The counterparts of these results hold in a constructive general lattice environment. Specifically, following [4], we call an inequality $\varphi \leq \psi$ *tame canonical* when $\mathbb{A} \models \varphi \leq \psi$ if and only if $\mathbb{A}^\delta \models \varphi^* \leq \psi^*$ for all μ-algebras $\mathbb{A}$ of the first kind. We generalize the *tame inductive* mu-inequalities of [4] to the LE-setting, and prove that they are tame canonical in a constructive meta-theory.

Of course, the usual notion of canonicity may also be applied to formulas with fixed-point binders, i.e., that $\mathbb{A} \models \varphi \leq \psi$ iff $\mathbb{A}^\delta \models \varphi \leq \psi$, where fixed points are interpreted in the standard way, e.g. least fixed points in $\mathbb{A}^\delta$ are calculated as the meet of all pre-fixed points *in* $\mathbb{A}^\delta$. A canonicity result of this kind can be proved by generalizing the class of *restricted inductive mu-inequalities* of [4] to the LE-setting, and showing that they are preserved under constructive canonical extensions of mu-algebras of the second kind.

Whenever a *tame* run of μ^*-ALBA succeeds on a mu-inequality $\varphi \leq \psi$, we have that $\varphi \leq \psi$ is *tame canonical*. Moreover, whenever a *proper* run succeeds on a mu-inequality $\alpha \leq \beta$, then $\alpha \leq \beta$ will be *canonical*. Finally, for every tame inductive mu-inequality (respectively, a restricted inductive mu-inequality), there exists a tame (respectively, proper) run of μ^*-ALBA which succeeds on that inequality.

4 Syntactic classes

In this section, we introduce some syntactically defined classes of mu-inequalities, the most general of which is the counterpart, in the language of normal LEs plus fixed points, of the recursive mu-inequalities introduced in [6]. In a constructive setting, the members of this class will all have correspondents in a first-order language plus fixed points. The remaining two classes of mu-inequalities defined in this section are subclasses of the class of recursive mu-inequalities, and are those for which the two canonicity results hold.

For any $\mathcal{L}_1$-sentence $\varphi(p_1, \ldots p_n)$, any order-type ε over n, and any $1 \leq i \leq n$, an *ε-critical node* in a signed generation tree of φ is a (leaf) node $+p_i$ with $\varepsilon_i = 1$, or $-p_i$ with $\varepsilon_i = \partial$. An *ε-critical branch* in the tree is a branch terminating in an ε-critical node. The intuition, which will be built upon later, is that variable occurrences corresponding to ε-critical nodes are *to be solved for, according to ε.*

In the signed generation tree of a $\mathcal{L}_1$-sentence $\varphi(p_1, \ldots p_n)$ a *live branch* is a branch ending in a (signed) propositional variable. In particular, all critical branches are live. A branch is not live iff it ends in a propositional constant ($\top$ or $\bot$) or in a fixed point variable.

For every $\mathcal{L}_1$-sentence $\varphi(p_1, \ldots p_n)$, and every order-type ε, we say that $+\varphi$ (resp. $-\varphi$) *agrees with* ε, and write $\varepsilon(+\varphi)$ (resp. $\varepsilon(-\varphi)$), if every leaf node in the signed generation tree of $+\varphi$ (resp. $-\varphi$) which is labelled with a propositional variable is ε-

Outer Skeleton (P_3)	Inner Skeleton (P_2)	PIA (P_1)
Δ-adjoints	Binders	Binders
$+ \quad \vee$	$+ \quad \mu$	$+ \quad \nu$
$- \quad \wedge$	$- \quad \nu$	$- \quad \mu$
SLR	SLA	SRA
$+ \quad f$	$+ \quad \vee \quad f \quad (n_f = 1)$	$+ \quad \wedge \quad g \quad (n_g = 1)$
$- \quad g$	$- \quad \wedge \quad g \quad (n_g = 1)$	$- \quad \vee \quad f \quad (n_f = 1)$
	SLR	SRR
	$+ \quad f \quad (n_f \geq 2)$	$+ \quad g \quad (n_g \geq 2)$
	$- \quad g \quad (n_g \geq 2)$	$- \quad f \quad (n_f \geq 2)$

Table 1
Skeleton and PIA nodes.

critical. We will also make use of the *sub-tree relation* $\gamma \prec \varphi$, which extends to signed generation trees, and we will write $\varepsilon(\gamma) \prec *\varphi$ to indicate that γ, regarded as a sub-(signed generation) tree of $*\varphi$, agrees with ε.

Definition 4.1 Nodes in signed generation trees will be called *skeleton nodes* and *PIA nodes* and further classified as Δ-adjoint, SLR, Binders, SLA, SRA or SRR, according to the specification given in Table 1.[4] Let $\varphi(p_1, \ldots, p_n)$ be a formula in the propositional variables $p_1, \ldots, p_n$, and let ε be an order-type on $\{1, \ldots, n\}$.

A branch in a signed generation tree $*\varphi$, for $* \in \{+, -\}$, ending in a propositional variable is an *ε-good branch* if, apart from the leaf, it is the concatenation of three paths P_1, P_2, and P_3, each of which may possibly be of length 0, such that P_1 is a path from the leaf consisting only of PIA-nodes, P_2 consists only of inner skeleton-nodes, and P_3 consists only of outer skeleton-nodes and, moreover, it satisfies conditions (GB1), (GB2) and (GB3), below.

(GB1) The formula corresponding to the uppermost node on P_1 is a sentence.

(GB2) For every SRR-node in P_1 of the form $h(\overline{\gamma}, \beta)$, where β is the coordinate where the branch lies, every γ in $\overline{\gamma}$ is a mu-sentence and $\varepsilon^\partial(\gamma) \prec *\varphi$ (i.e., each γ contains no variable occurrences to be solved for — see above).

(GB3) For every SLR-node in P_2 of the form $h(\overline{\gamma}, \beta)$, where β is the coordinate where the branch lies, every γ in $\overline{\gamma}$ is a mu-sentence and $\varepsilon^\partial(\gamma) \prec *\varphi$ (see above for this notation).

Our main interest is in ε-good branches satisfying some of the additional properties in the following definition.

Definition 4.2 Let $\varphi(p_1, \ldots, p_n)$ be a formula, ε be an order-type on $\{1, \ldots, n\}$, and $<_\Omega$ a strict partial order on the variables $p_1, \ldots p_n$. An ε-good branch may satisfy one or more of the following properties:

[4] The abbreviations SLR, SLA, SRA and SRR stand for syntactically left residual, left adjoint, right adjoint and right residual, respectively. Nodes are thus classified according to the order-theoretic properties of their interpretations, see [8,11] for further discussion on methodology and nomenclature.

(NB-PIA) P_1 contains no fixed point binders.

(NL) For every SLR-node in P_2 of the form $h(\overline{\gamma}, \beta)$, where β is the coordinate where the branch lies, the signed generation tree of each γ contains no live branches.

(Ω-CONF) For every SRR-node in P_1 of the form $h(\overline{\gamma}, \beta)$, where β is the coordinate where the branch lies, $p_j <_\Omega p_i$ for every p_j occurring in γ, where p_i is the propositional variable labelling the leaf of the branch.

Definition 4.3 For any order-type ε and strict partial order $<_\Omega$ on the variables $p_1, \dots p_n$, the signed generation tree $*\varphi$, $* \in \{-, +\}$, of a term $\varphi(p_1, \dots p_n)$ is called

(i) ε-*recursive* if every ε-critical branch is ε-good.

(ii) (Ω, ε)-*inductive* it is ε-recursive and every ε-critical branch satisfies (Ω-CONF).

(iii) *restricted* (Ω, ε)-*inductive* if it is (Ω, ε)-inductive and
   (a) every ε-critical branch satisfies (NB-PIA) and (NL),
   (b) every occurrence of a binder is on an ε-critical branch.

(iv) *tame* (Ω, ε)-*inductive* if it is (Ω, ε)-inductive and
   (a) $\Omega = \varnothing$,
   (b) no binder occurs on any ε-critical branch,
   (c) the only nodes involving binders which are allowed to occur are $+\nu$ and $-\mu$.

An inequality $\varphi \leq \psi$ is ε-recursive (resp., (Ω, ε)-inductive, restricted (Ω, ε)-inductive, tame (Ω, ε)-inductive) if $+\varphi$ and $-\psi$ are both ε-recursive (resp., (Ω, ε)-inductive, restricted (Ω, ε)-inductive, tame (Ω, ε)-inductive).

An inequality $\varphi \leq \psi$ is recursive (resp., inductive, restricted inductive, tame inductive) if $\varphi \leq \psi$ is ε-recursive (resp., (Ω, ε)-inductive, restricted (Ω, ε)-inductive, tame (Ω, ε)-inductive) for some strict partial order Ω and order-type ε.

The corresponding classes of inequalities will be referred to as the *recursive* (resp., *inductive, restricted inductive, tame inductive*) *mu-inequalities*, or the *recursive* (resp., *inductive, restricted inductive, tame inductive*) *mu-formulas*, if the inequality signs have been replaced with implications.

5 Constructive μ^*-ALBA

In this section, we introduce the constructive and general lattice version of the algorithm μ^*-ALBA introduced in [4], which in its turn is a restricted version of the algorithm μ-ALBA, introduced in [6] to calculate first-order correspondents of recursive inequalities from bi-intuitionistic modal mu-calculus. Constructive μ^*-ALBA is the fundamental tool to prove the canonicity results via the argument discussed in Section 3. As usual, the goal of constructive μ^*-ALBA is to *eliminate propositional variables* from inequalities, while maintaining admissible validity. The purpose of this is to make the transition from *admissible validity* to *validity* in the argument for canonicity. Below, we outline its general strategy.

Constructive μ^*-ALBA, from now on abbreviated as ALBA, takes an $\mathcal{L}_1$-inequality $\varphi \leq \psi$ as input, and proceeds in three stages. The first stage preprocesses by eliminating all uniformly occurring propositional variables, applying distribution rules for $f \in \mathcal{F}$ and $g \in \mathcal{G}$ and splitting rules exhaustively, and converting all occurrences

of $\mu X.\varphi(X)$ to $\mu^* X.\varphi(X)$ and all occurrences of $\nu X.\psi(X)$ to $\nu^* X.\psi(X)$. We emphasize that this step is required in both tame *and* proper runs of ALBA (see Appendix A). The preprocessing produces a finite set of inequalities, $\varphi'_i \leq \psi'_i$, $1 \leq i \leq n$.

Now ALBA forms the *initial quasi-inequalities* $\&\, S_i \Rightarrow \mathsf{Ineq}_i$, compactly represented as tuples (S_i, Ineq_i) referred to as *systems*, with each S_i initialized to the empty set and Ineq_i initialized to $\varphi'_i \leq \psi'_i$.

The second stage (called the reduction stage) transforms S_i and Ineq_i through the application of transformation rules (see Appendix A). The aim is to eliminate all propositional variables from S_i and Ineq_i in favour of nominals and co-nominals. A system for which this has been done will be called *pure* or *purified*. The actual eliminations are effected through the Ackermann rules, while the other rules are used to bring S_i and Ineq_i into the appropriate shape which make these applications possible. Once all propositional variables have been eliminated, this phase terminates and returns the pure quasi-inequalities $\&\, S_i \Rightarrow \mathsf{Ineq}_i$.

The third stage either reports failure if some system could not be purified, or else returns the conjunction of the pure quasi-inequalities $\&\, S_i \Rightarrow \mathsf{Ineq}_i$, which we denote by $\mathsf{ALBA}(\varphi \leq \psi)$.

A more complete outline of each of the three stages will be given in Appendix A.

6 Main results

In this section we state the main results of the paper, namely that all restricted inductive inequalities are constructively canonical and that all tame inductive inequalities are constrictively tame canonical. The proof strategy is the same in both cases: one first proves, by means of a 'U-shaped' argument as discussed in Section 3, that successful runs of constructive ALBA satisfying certain conditions guarantee these types of canonicity. Then it is shown that all members of the two classes of inequalities are successfully reducible by means of runs respectively satisfying these properties. The main canonicity results then follow as corollaries. Because of space limitations we do not include the proofs.

A proof combining insights from the proofs of [4, Theorems 9.9 and 9.4] and [9, Theorem 7.1] suffices to establish the following two propositions:

Proposition 6.1 *Let $\mathbb{A}$ be a mu-algebra of the second kind and let $\varphi \leq \psi$ be an $\mathcal{L}_1$-inequality on which a proper and pivotal run of μ^*-ALBA succeeds. If $\mathbb{A} \models \varphi \leq \psi$ then $\mathbb{A}^\delta \models \varphi \leq \psi$.*

Proposition 6.2 *All $\mathcal{L}_1$-inequalities on which a tame and pivotal run of constructive μ^*-ALBA succeeds are constructively tame canonical.*

Again, generalizing to the non-distributive environment and amalgamating the strategies from [4, Section 10] and [9, Section 6] yields the next proposition:

Proposition 6.3 *Constructive μ^*-ALBA succeeds on all restricted inductive $\mathcal{L}$-inequalities by means of proper and pivotal runs.*

Proposition 6.4 *Constructive μ^*-ALBA succeeds on all tame inductive $\mathcal{L}$-inequalities by means of tame and pivotal runs.*

The canonicity of restricted inductive $\mathcal{L}$-inequalities now follows as a corollary of

Propositions 6.1 and 6.3.

Theorem 6.5 *All restricted inductive $\mathcal{L}$-inequalities are constructively canonical over mu-algebras of the second kind.*

Similarly, the tame canonicity of all tame inductive inequalities follows from Propositions 6.2 and 6.4.

Theorem 6.6 *All tame inductive $\mathcal{L}$-inequalities are constructively canonical over mu-algebras of the first kind.*

References

[1] Ambler, S., M. Kwiatkowska and N. Measor, *Duality and the completeness of the modal μ-calculus*, Theoretical Computer Science **151** (1995), pp. 3–27.

[2] Baelde, D., *Least and greatest fixed points in linear logic*, ACM Transactions on Computational Logic (TOCL) **13** (2012), p. 2.

[3] Bezhanishvili, N. and I. Hodkinson, *Sahlqvist theorem for modal fixed point logic*, Theoretical Computer Science **424** (2012), pp. 1–19.

[4] Conradie, W. and A. Craig, *Canonicity results for mu-calculi: An algorithmic approach*, Journal of Logic and Computation (forthcoming).
URL `http://arxiv.org/abs/1408.6367`

[5] Conradie, W., A. Craig, A. Palmigiano and Z. Zhao, *Constructive canonicity for lattice-based fixed point logics*, ArXiv preprint .

[6] Conradie, W., Y. Fomatati, A. Palmigiano and S. Sourabh, *Algorithmic correspondence for intuitionistic modal mu-calculus*, Theoretical Computer Science **564** (2015), pp. 30–62.

[7] Conradie, W., S. Frittella, A. Palmigiano, M. Piazzai, A. Tzimoulis and N. Wijnberg, *Reasoning about categories with generalized Kripke frames*, in preparation .

[8] Conradie, W., S. Ghilardi and A. Palmigiano, *Unified correspondence*, in: A. Baltag and S. Smets, editors, *Johan van Benthem on Logic and Information Dynamics*, Outstanding Contributions to Logic **5**, Springer International Publishing, 2014 pp. 933–975.

[9] Conradie, W. and A. Palmigiano, *Constructive canonicity of inductive inequalities* Submitted.

[10] Conradie, W. and A. Palmigiano, *Algorithmic correspondence and canonicity for distributive modal logic*, Annals of Pure and Applied Logic **163** (2012), pp. 338 – 376.

[11] Conradie, W. and A. Palmigiano, *Algorithmic correspondence and canonicity for non-distributive logics*, Journal of Logic and Computation (forthcoming).

[12] Conradie, W., A. Palmigiano and S. Sourabh, *Algorithmic modal correspondence: Sahlqvist and beyond* Submitted.

[13] Conradie, W., A. Palmigiano, S. Sourabh and Z. Zhao, *Canonicity and relativized canonicity via pseudo-correspondence: an application of ALBA* Submitted.

[14] Conradie, W., A. Palmigiano and Z. Zhao, *Sahlqvist via translation* Submitted.

[15] Conradie, W. and C. Robinson, *On Sahlqvist theory for hybrid logic*, Journal of Logic and Computation (2015).

[16] Dunn, J. M., M. Gehrke and A. Palmigiano, *Canonical extensions and relational completeness of some substructural logics*, Journal Symbolic Logic **70** (2005), pp. 713–740.

[17] Frittella, S., A. Palmigiano and L. Santocanale, *Dual characterizations for finite lattices via correspondence theory for monotone modal logic*, Journal of Logic and Computation (forthcoming).

[18] Gavazzo, F., *Investigations into linear logic with fixed-point operators* (2015), iLLC MoL Thesis.

[19] Gehrke, M. and J. Harding, *Bounded lattice expansions*, Journal of Algebra **238** (2001), pp. 345–371.

[20] Ghilardi, S. and G. Meloni, *Constructive canonicity in non-classical logics*, Annals of Pure and Applied Logic **86** (1997), pp. 1–32.

[21] Greco, G., M. Ma, A. Palmigiano, A. Tzimoulis and Z. Zhao, *Unified correspondence as a proof-theoretic tool*, Journal of Logic and Computation (forthcoming).

[22] Hartonas, C., *Duality for modal μ-logics*, Theoretical Computer Science **202** (1998), pp. 193–222.

[23] Ma, M. and Z. Zhao, *Unified correspondence and proof theory for strict implication*, Journal of Logic and Computation (forthcoming).

[24] Palmigiano, A., S. Sourabh and Z. Zhao, *Jónsson-style canonicity for ALBA-inequalities*, Journal of Logic and Computation (2015).

[25] Palmigiano, A., S. Sourabh and Z. Zhao, *Sahlqvist theory for impossible worlds*, Journal of Logic and Computation (forthcoming).

[26] Santocanale, L., *Completions of μ-algebras*, in: *Logic in Computer Science, 2005. LICS 2005. Proceedings. 20th Annual IEEE Symposium on*, IEEE, 2005, pp. 219–228.

[27] Suzuki, T., *Canonicity results of substructural and lattice-based logics*, The Review of Symbolic Logic **4** (2011), pp. 1–42.

Appendix

A Stages and rules of μ^*-ALBA

A.1 Stage 1: Preprocessing and initialization

ALBA receives an $\mathcal{L}_1$-inequality $\varphi \leq \psi$ as input. It applies the following **rules for elimination of monotone variables** to $\varphi \leq \psi$ exhaustively, in order to eliminate any propositional variables which occur uniformly:

$$\frac{\alpha(p) \leq \beta(p)}{\alpha(\bot) \leq \beta(\bot)} \qquad \frac{\gamma(p) \leq \delta(p)}{\gamma(\top) \leq \delta(\top)}$$

for $\alpha(p) \leq \beta(p)$ positive and $\gamma(p) \leq \delta(p)$ negative in p, respectively.[5]

Next, ALBA exhaustively distributes $f \in \mathcal{F}$ over $+\vee$, and $g \in \mathcal{F}$ over $-\wedge$, so as to bring occurrences of $+\vee$ and $-\wedge$ to the surface wherever this is possible, and then eliminate them via exhaustive applications of *splitting* rules.

Splitting-rules.

$$\frac{\alpha \leq \beta \wedge \gamma}{\alpha \leq \beta \quad \alpha \leq \gamma} \qquad \frac{\alpha \vee \beta \leq \gamma}{\alpha \leq \gamma \quad \beta \leq \gamma}$$

This gives rise to a set of inequalities $\{\varphi'_i \leq \psi'_i \mid 1 \leq i \leq n\}$. For each of them, ALBA converts all occurrences of $\mu X.\varphi(X)$ to $\mu^* X.\varphi(X)$ and all occurrences of $\nu X.\psi(X)$ to $\nu^* X.\psi(X)$, and forms the *initial quasi-inequality* $\&S_i \Rightarrow \mathsf{Ineq}_i$, compactly represented as a tuple (S_i, Ineq_i) referred as *initial system*, with each S_i initialized to the empty set and Ineq_i initialized to $\varphi'_i \leq \psi'_i$. Each initial system is passed separately to stage 2, described below, where we will suppress indices i.

A.2 Stage 2: Reduction and elimination

The aim of this stage is to eliminate all occurring propositional variables from a given system (S, Ineq). This is done by means of the following *approximation rules*, *residuation rules*, *splitting rules*, and *Ackermann rules*, collectively called *reduction rules*. The terms and inequalities in this subsection are from $\mathcal{L}^+_*$.

[5] A term φ is *positive* (*negative*) in a variable p if in the generation tree $+\varphi$ all p-nodes are signed $+$ ($-$). An inequality $\varphi \leq \psi$ is *positive* (*negative*) in p if φ is negative (positive) in p and ψ is positive (negative) in p.

Approximation rules. There are four approximation rules. Each of these rules functions by simplifying Ineq and adding an inequality to S. We write $\alpha(!x)$ to indicate that the placeholder variable x has a unique occurrence in formula α.

Left-positive approximation rule.

$$\frac{(S,\ \ \varphi'(\gamma/!x) \leq \psi)}{(S \cup \{\mathbf{j} \leq \gamma\},\ \ \varphi'(\mathbf{j}/!x) \leq \psi)}\ (L^+A)$$

with $+x \prec +\varphi'(!x)$, the branch of $+\varphi'(!x)$ starting at $+x$ subject to the restrictions detailed below, γ belonging to the smaller language $\mathcal{L}_*$ and $\mathbf{j}$ being the first nominal variable not occurring in S or $\varphi'(\gamma/!x) \leq \psi$.

Left-negative approximation rule.

$$\frac{(S,\ \ \varphi'(\gamma/!x) \leq \psi)}{(S \cup \{\gamma \leq \mathbf{m}\},\ \ \varphi'(\mathbf{m}/!x) \leq \psi)}\ (L^-A)$$

with $-x \prec +\varphi'(!x)$, the branch of $+\varphi'(!x)$ starting at $-x$ subject to the restrictions detailed below, γ belonging to the smaller language $\mathcal{L}_*$ and $\mathbf{m}$ being the first conominal not occurring in S or $\varphi'(\gamma/!x) \leq \psi$.

Right-positive approximation rule.

$$\frac{(S,\ \ \varphi \leq \psi'(\gamma/!x))}{(S \cup \{\mathbf{j} \leq \gamma\},\ \ \varphi \leq \psi'(\mathbf{j}/!x))}\ (R^+A)$$

with $+x \prec -\psi'(!x)$, the branch of $-\psi'(!x)$ starting at $+x$ subject to the restrictions detailed below, γ belonging to the smaller language $\mathcal{L}_*$ and $\mathbf{j}$ being the first nominal not occurring in S or $\varphi \leq \psi'(\gamma/!x)$.

Right-negative approximation rule.

$$\frac{(S,\ \ \varphi \leq \psi'(\gamma/!x))}{(S \cup \{\gamma \leq \mathbf{m}\},\ \ \varphi \leq \psi'(\mathbf{m}/!x))}\ (R^-A)$$

with $-x \prec -\psi'(!x)$, the branch of $-\psi'(!x)$ starting at $-x$ subject to the restrictions detailed below, γ belonging to the smaller language $\mathcal{L}_*$ and $\mathbf{m}$ being the first conominal not occurring in S or $\varphi \leq \psi'(\gamma/!x)$).

The restrictions on φ' and ψ' in the approximation rules above are formulated in terms of the following:

Definition A.1 For any mu-algebra $\mathbb{C}$ and order-type τ, a join $\bigvee S$ in $(\mathbb{C}^\delta)^\tau$ is called $\mathbb{C}^\tau$-*targeted* if $\bigvee S \in \mathbb{C}^\tau$. A map $f : (\mathbb{C}^\delta)^\tau \to \mathbb{C}^\delta$ *preserves* $\mathbb{C}^\tau$-*targeted joins* if $f(\bigvee S) = \bigvee_{s \in S} f(s)$ for every $S \subseteq (\mathbb{C}^\delta)^\tau$ such that $\bigvee S$ is $\mathbb{C}^\tau$-targeted. *Targeted meets* and their preservation are defined order-dually.

Let us now list the requirements on φ' and ψ':

(i) $\varphi', \psi', \in \mathcal{L}_*$;

(ii) the branches of φ' and ψ' starting at x going up to the root consist only of Skeleton nodes. [6]

[6] The purpose of this restriction is to enforce preservation of non-empty joins by the term function $\varphi'^{\mathbb{C}}$. The soundness of the rule is founded upon this and approximation of the argument γ as the join of all closed elements below it. In the non-constructive setting of [11] the same strategy is followed, except that

(iii) for every node of the form $+\mu X.\psi(\overline{x}, X)$ or of the form $-\nu X.\varphi(\overline{x}, X)$ in such branches, which is not in the scope of another binder, all propositional variables and free fixed point variables in $\psi(\overline{x}, X)$ and $\varphi(\overline{x}, X)$ must be among $\overline{x}$ and X; moreover,
   (a) the associated term function $\psi(\overline{x}, X) : (\mathbb{C}^\delta)^\tau \times \mathbb{C}^\delta \to \mathbb{C}^\delta$ preserves $(\mathbb{C}^\tau \times \mathbb{C})$-targeted joins for all μ-algebras $\mathbb{C}$ of the second kind; moreover $\psi(\overline{x}, X)$ is required to be positive (negative) in x_i if $\tau_i = 1$ ($\tau_i = \partial$), i.e. $\psi(\overline{x}, X)$ must be τ-positive in $\overline{x}$;
   (b) the associated term function $\varphi(\overline{x}, X) : (\mathbb{C}^\delta)^\tau \times \mathbb{C}^\delta \to \mathbb{C}^\delta$ preserves $(\mathbb{C}^\tau \times \mathbb{C})$-targeted meets for all $\mathcal{L}_{\mathrm{LE}}$-algebras $\mathbb{C}$ of the second kind; moreover $\varphi(\overline{x}, X)$ is required to be positive (negative) in x_i if $\tau_i = 1$ ($\tau_i = \partial$), i.e. $\varphi(\overline{x}, X)$ must be τ-positive in $\overline{x}$.

Remark A.2 (i) The approximation rules above, as stated, are sound both under admissible and under arbitrary assignments. However, their liberal application gives rise to topological complications in the proof of canonicity. Therefore, we will restrict the applications of approximation rules to nodes $!x$ giving rise to *maximal* skeleton branches.Such applications will be called *pivotal*. Also, executions of ALBA in which approximation rules are applied only pivotally will be referred to as *pivotal*.

(ii) In [4], approximation rules were formulated specifically for formulas having a fixed point binder as main connective. These rules had a substantially more cumbersome formulation than the one given above, which, modulo the restrictions about the preservation of targeted joins and meets, follows verbatim the approximation rules of [11]. Moreover, the approximation rules [4] could give rise to the splitting of the quasi-inequality into a set of quasi-inequalities, which is not the case of the present setting. This is thanks to the fact that nominals and co-nominals are not interpreted as completely join-primes (resp. meet-primes), but as closed and open elements, and this notion is compatible with products (cf. Section 2, page 5).

Residuation rules. These rules operate on the inequalities in S, by rewriting a chosen inequality in S into another inequality. For every $f \in \mathcal{F}$ and $g \in \mathcal{G}$, and any $1 \leq i \leq n_f$ and $1 \leq j \leq n_g$,

$$\frac{f(\varphi_1, \ldots, \varphi_i, \ldots, \varphi_{n_f}) \leq \psi}{\varphi_i \leq f_i^\sharp(\varphi_1, \ldots, \psi, \ldots, \varphi_{n_f})}\, \varepsilon_f(i) = 1 \qquad \frac{f(\varphi_1, \ldots, \varphi_i, \ldots, \varphi_{n_f}) \leq \psi}{f_i^\sharp(\varphi_1, \ldots, \psi, \ldots, \varphi_{n_f}) \leq \varphi_i}\, \varepsilon_f(i) = \partial$$

$$\frac{\psi \leq g(\varphi_1, \ldots, \varphi_i, \ldots, \varphi_{n_g})}{g_i^\flat(\varphi_1, \ldots, \psi, \ldots, \varphi_{n_g}) \leq \varphi_i}\, \varepsilon_g(i) = 1 \qquad \frac{\psi \leq g(\varphi_1, \ldots, \varphi_i, \ldots, \varphi_{n_g})}{\varphi_i \leq g_i^\flat(\varphi_1, \ldots, \psi, \ldots, \varphi_{n_g})}\, \varepsilon_g(i) = \partial$$

---

the approximation is done by means of completely join-irreducibles. Since this can give rise to empty sets of approximants and hence empty joins, $+\vee$ is excluded in the analogous approximation rule in [11], as the join does not preserve empty joins coordinate-wise. In the present setting, the set of closed approximants is never empty, and hence this restriction may be dropped. Similar considerations apply to $-\wedge$.

In a given system, each of these rules replaces an instance of the upper inequality with the corresponding instances of the two lower inequalities.

Ackermann rules The Ackermann rules are used for the crucial task of eliminating propositional variables from quasi-inequalities.

$$\frac{\exists p[\&_{i=1}^{n}\alpha_i \leq p \ \& \ \&_{j=1}^{m}\beta_j(p) \leq \gamma_j(p)]}{\&_{j=1}^{m}\beta_j(\vee_{i=1}^{n}\alpha_i/p) \leq \gamma_j(\vee_{i=1}^{n}\alpha_i/p)}\ (\mathrm{RA})$$

subject to the restrictions that the α_i are p-free and syntactically closed, the β_j are positive in p and syntactically closed, while the γ_j are negative in p and syntactically open (cf. subsection below).

$$\frac{\exists p[\&_{i=1}^{n} p \leq \alpha_i \ \& \ \&_{j=1}^{m}\gamma_j(p) \leq \beta_j(p)]}{\&_{j=1}^{m}\gamma_j(\wedge_{i=1}^{n}\alpha_i/p) \leq \beta_j(\wedge_{i=1}^{n}\alpha_i/p)}\ (\mathrm{LA})$$

subject to the restrictions that the α_i are p-free and syntactically open, the β_j are positive in p and syntactically open, while the γ_j are negative in p and syntactically closed.

A.2.1 Syntactically (almost) open and closed formulas

As mentioned in the introduction, when formulas from the extended language $\mathcal{L}^+$ are interpreted in $\mathbb{A}^\delta$, an assignment V will have $V(\mathsf{NOM}) \subseteq K(\mathbb{A}^\delta)$ and $V(\mathsf{CNOM}) \subseteq O(\mathbb{A}^\delta)$. For any assignment V, the assignment V' is a p-variant (or, a $\mathbf{j}$-variant or $\mathbf{m}$-variant) of V if V' agrees with V on all elements of $\mathsf{PROP} \cup \mathsf{NOM} \cup \mathsf{CNOM}$ except possibly at p (respectively, at $\mathbf{j}$ or $\mathbf{m}$). If so, we write $V' \sim_p V$ (respectively, $V' \sim_{\mathbf{j}} V$ or $V' \sim_{\mathbf{m}} V$). From this point on, we will use V to denote both the assignment $V : \mathsf{PROP} \cup \mathsf{NOM} \cup \mathsf{CNOM} \to \mathbb{A}^\delta$ and its unique homomorphic extension $V : \mathcal{L}^+ \to \mathbb{A}^\delta$.

In the following definition we will use $f \in \mathcal{F}$ and $g \in \mathcal{G}$ to denote connectives of the original signature, and $h \in \mathcal{F}^+ \setminus \mathcal{F}$ and $k \in \mathcal{G}^+ \setminus \mathcal{G}$ to denote connectives of the expanded ('tense') language. To simplify notation, we will disregard the actual order of the coordinates, but keep track of their polarity. So, for instance we will write $f(\overline{\psi}, \overline{\varphi})$ and $k(\overline{\varphi}, \overline{\psi})$, where in both cases the coordinates are divided in two possibly empty arrays, the first (resp. second) of which contains the positive (resp. negative) coordinates.

Definition A.3 The *syntactically open formulas* φ and *syntactically closed formulas* ψ are defined by simultaneous mutual recursion as follows:

$$\varphi ::= \bot \mid \top \mid p \mid \mathbf{m} \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \vee \varphi_2 \mid g(\overline{\varphi}, \overline{\psi}) \mid f(\overline{\varphi}, \overline{\psi}) \mid k(\overline{\varphi}, \overline{\psi}) \mid \nu^* X.\varphi$$
$$\psi ::= \bot \mid \top \mid p \mid \mathbf{i} \mid \psi_1 \wedge \psi_2 \mid \psi_1 \vee \psi_2 \mid f(\overline{\psi}, \overline{\varphi}) \mid g(\overline{\psi}, \overline{\varphi}) \mid h(\overline{\psi}, \overline{\varphi}) \mid \mu^* X.\psi$$

where $p \in \mathsf{PROP}$, $\mathbf{i} \in \mathsf{NOM}$, and $\mathbf{m} \in \mathsf{CNOM}$.

The *syntactically almost open formulas* and *syntactically almost closed formulas* are defined by adding $\mu^* X.\varphi$ (resp. $\nu^* X.\psi$) to the recursive definition of φ (resp. ψ).

Informally, an $\mathcal{L}^+_*$-term is *syntactically almost open* (resp. *syntactically almost closed*) if, in it, all occurrences of nominals and $h \in \mathcal{F}^+ \setminus \mathcal{F}$ are negative (resp. positive), while all occurrences of co-nominals and $k \in \mathcal{G}^+ \setminus \mathcal{G}$ are positive (resp. negative).

If, in addition, all occurrences of μ^* are negative (resp. positive) and all occurrences of ν^* positive (resp. negative), the term is *syntactically open* (resp. *syntactically closed*).

A.3 Stage 3: Success, failure and output

If stage 2 succeeded in eliminating all propositional variables from each system, the algorithm returns the conjunction of these purified quasi-inequalities, denoted by $\mathrm{ALBA}(\varphi \leq \psi)$. Otherwise, the algorithm reports failure and terminates.

A *tame run* of μ^*-ALBA is one during which the approximation rules are applied only to formulas $\varphi'(\gamma/!x)$ and $\psi'(\gamma/!x)$ such that no fixed point binder occurs in the branch from x to the root of φ' and ψ'. By contrast, a *proper run* of μ^*-ALBA is one during which *all* occurrences of fixed point binders lie along some branch ending with a subterm γ which the application of an approximation rule extracts. We say that a run of the algorithm μ^*-ALBA *succeeds* if all propositional variables are eliminated from the input inequality, $\varphi \leq \psi$, and denote the resulting set of pure quasi-inequalities by $\mathrm{ALBA}(\varphi^* \leq \psi^*)$. An inequality on which some run of μ^*-ALBA succeeds is called a *μ^*-ALBA inequality*.